

DEVELOPMENT OF AN ITS DATA MANAGEMENT SYSTEM

by

Robert E. Brydia
Assistant Research Scientist
E-mail: reb@myriad.net

Shawn M. Turner
Assistant Research Scientist
E-mail: shawn-turner@tamu.edu

William L. Eisele
Assistant Research Scientist
E-mail: bill-eisle@tamu.edu

of the

Texas Transportation Institute
The Texas A&M University System
College Station, Texas 77843-3135
Phone (409) 845-1536
Fax (409) 845-6001

and

Jyh C. Liu
Assistant Professor
Department of Computer Science
Texas A&M University
E-mail: jcliu@cs.tamu.edu

Paper Prepared for Presentation and Publication

at the

Transportation Research Board
77th Annual Meeting
Washington, D.C.

January 1998

Word Count: 6,286 words (5,286 words + 1 figure + 3 tables)

ABSTRACT

The intelligent transportation system (ITS) components being deployed in U.S. urban areas are producing vast amounts of data. These ITS data are often used for real-time operations and then discarded. Few transportation management centers (TMCs) have any mechanism for sharing the data resources among other transportation groups or agencies within the same jurisdiction. Meanwhile, transportation analysts and researchers often struggle to obtain accurate, reliable data about existing transportation performance and patterns. This paper presents the development of an ITS data management system (referred to as ITS DataLink) that is used to store, access, analyze, and present data from the TransGuide center in San Antonio, Texas. Data outputs are both tabular and graphical in nature. There are no user costs associated with system except for an Internet connection.

INTRODUCTION

The intelligent transportation system (ITS) components being deployed in U.S. urban areas are producing vast amounts of data. Vehicle detectors collect information about vehicle speeds, volumes, and lane occupancies. Probe vehicle systems enable the collection of travel times and average speeds along roadway sections. These ITS data are often used for real-time operations and then discarded. Some transportation management centers (TMCs) may save the data, but few have any mechanism for sharing the data resources among other transportation groups or agencies within the same jurisdiction. Even if the data are saved, many people are faced with issues related to transforming the vast amounts of data into useful information for adjusting operating strategies, evaluating system performance, or making decisions about future transportation investments.

Many transportation analysts and researchers struggle to obtain accurate, reliable data about existing transportation performance and patterns. Models rely heavily on existing conditions for calibration purposes, and decision-makers rely on models and the existing transportation performance to make decisions about transportation investments. The importance of accurate, reliable data in transportation analyses is paramount to sound decisions in planning, designing, operating, and maintaining the transportation system.

This paper presents the development of an ITS data management system (referred to as ITS DataLink) that is used to store, access, analyze, and present data from the TransGuide center in San Antonio, Texas.

REVIEW OF DATA MANAGEMENT TECHNIQUES

As part of the research, efforts were undertaken to examine how TMCs around the country were utilizing the vast amounts of data that are available through ITS components. Some of the many questions associated with the issue of ITS data utilization include:

- How are the data being collected [e.g., inductance loops, Automatic Vehicle Identification (AVI)];
- What data are being collected (e.g., volume, speed, occupancy, classification);
- To what extent are the data are archived and aggregated (e.g., 5 minute, 15 minute);
- Who are the most common users of the data (i.e., typical agencies requesting information);
- For what purposes are the data being used (i.e. planning, design, or operational purposes or some combination) and;

- Are there any additional data management considerations (i.e., privacy issues and/or established policies for data sharing).

A survey was developed and conducted by phone to determine what data management strategies are in existence and what opportunities may exist for utilizing ITS data. Table 1 lists the 15 TMCs listed that were interviewed in the study.

Based on the survey results, although the uses and users of ITS data appear to be consistent across the board, the aggregation and storage practices are not. Fully 20 percent of the TMCs are not saving any ITS data. Sixty percent of the TMCs are saving data at an aggregation level of one minute or less. The remaining 20 percent store their data with an aggregation level between 5 and 15 minutes.

SYSTEM GOALS

In any system design, it is necessary to define the desired goals and then structure the system accordingly. In the case of the prototype ITS DataLink system the desired goal was to build a system that could quickly and easily allow a user to determine system performance parameters that would apply to a variety of applications. The San Antonio TransGuide system was selected as the TMCs for which the system would be developed. In large part, this was due to the ample amount of data available from the TMCs as TransGuide stores their ITS data with no aggregation at all. This allowed the researchers to have complete system data every 20 seconds.

Additional goals of the prototype system included the questions below:

- What are the needs and potential uses of ITS data beyond real-time applications?
- What performance measures are needed for these applications, and can the available ITS data be used to calculate these performance measures?
- What are the issues associated with the storage, aggregation and analysis of ITS data?
- What computer hardware and software is required for an ITS data management system?

PERFORMANCE MEASURES FOR ITS DATA

Performance measures are used to quantify the system status. These measures are spatially distributed, i.e. different measures apply to different geographical summaries (point, link, and corridor) of the baseline data. Table 2 lists the performance measures considered for the prototype system.

Data collected by ITS components are based upon the needs and specific functions of the given system component. For example, inductance loops and other point detection devices are designed

to provide information about traffic conditions on a frequent basis. As such, inductance loops and other point detection devices typically report lane-by-lane speeds, volumes, and occupancies every 20 to 30 seconds. This detailed information about traffic conditions is then used in “real-time” for a variety of applications that includes, but is not limited to:

- Verification of and response to incidents;
- Operation of ramp metering strategies;
- Operation of changeable message and lane assignment signs;
- Operation of traffic-adaptive signal control strategies; and,
- Provision of traffic condition information to travelers.

Each specific application of ITS data has certain requirements in terms of level of detail and format. For the previous example, the real-time operational applications require detailed information (e.g., typically lane-by-lane detectors every 0.8 km or 0.5 mile) that is updated on a frequent basis (e.g., every 20 to 30 seconds). Secondary uses or applications of ITS data related to transportation planning, programming, or evaluation typically require less detailed information (e.g., corridor or system) for extended periods of time (e.g., monthly or annual averages). Thus it is vitally important to consider the needs and potential uses of ITS data in developing an ITS data management system, as they control the level of detail for data storage and analyses.

Based on the survey results and additional information, the following important points about ITS data should be noted:

- Different uses of ITS data require different levels of detail;
- Planning applications commonly require historical data over extended sections of roadway and periods of time;
- Design and operational applications commonly require detailed data for shorter sections and roadway and small intervals of time; and,
- Evaluations require a range of detail levels for both time and space.

Types of Performance Measures

Performance measures related to the efficiency and quality of service that can be calculated or estimated using readily available ITS data include:

Point Measures

Point measures are associated with a particular spot location on the transportation system. They typically include the following:

- spot speed (kilometers per hour); and,
- person volume or throughput (persons per hour).

Link-Based Measures

Link-Based Measures are usually associated with a small sub-portion of a person trip or transportation facility. Typical link performance measures are:

- travel time (seconds or minutes);
- average speed (kilometers per hour);
- person volume or throughput (persons per hour);
- person-movement speed (person-kilometers per hour); and,
- person delay (person-hours).

Corridor or System Measures

Corridor or system measures are the largest aggregation level of ITS data and are usually associated with a large sub-portion of a person trip or transportation facility. Typical measures are:

- average person speed (person-kilometers per hour);
- total person delay (person-hours);
- person-miles of travel in congestion (total and percent);
- person-hours of travel in congestion (total and percent);
- corridor mobility index; and,
- roadway congestion index.

ISSUES ASSOCIATED WITH ITS DATA

While the calculation of performance measures is typically straightforward, several issues do arise in retaining, managing, and analyzing ITS data for any purpose.

- Data storage--because of the enormous volumes of ITS data being collected, innovative storage and/or aggregation strategies are necessary to keep costs to a reasonable level and within the reach of a typical agency;
- Data Base Construction--the potentially large ITS data bases cannot be built with most traditional desktop computer-based spreadsheet or data base applications;
- Access to Data--the ITS data bases cannot be accessed without knowing specialized applications or query languages to interact with the data base engine;
- Data versus information--the sheer size of ITS data bases may make it difficult to convert megabytes of data to smaller, easy-to-understand information; and,
- Privacy--the possibility of having information recorded in ITS efforts that pinpoints where an individual was and when raises some privacy concerns and issues.

Many transportation management centers do not retain ITS data because of one or more of these issues, or because they haven't recognized the usefulness of the data. These issues are discussed in the following sections.

Data Storage

The enormous volumes of ITS data being collected clearly require innovative storage and/or aggregation strategies. As of August 1997, the San Antonio loop detector data is being archived as compressed ASCII-text files on a computer workstation. Each day of compressed data requires approximately 12 megabytes (MB) of storage space (uncompressed storage requires about 120 MB per day). At this rate of storage, a month of compressed data requires 360 MB (3.6 gigabytes (GB) uncompressed) and a full year of compressed data requires 4.4 GB (44 GB uncompressed). Also, these storage requirements are only for 42 km (26 mi) of freeway in Phase One of TransGuide. Phase Two will more than double the amount of freeway coverage and will necessarily increase the data storage space by roughly the same multiplier. Although data base or other formats may slightly decrease the storage space required, it is unlikely that current desktop database applications managing tens of gigabytes of data can be used as an effective data storage and management solutions. In addition, the choice of the operating system for the application is important, as some operating systems are better designed for data intensive tasks than others.

Many large businesses and corporations use large "enterprise-class" relational data bases on computer workstations to manage data requiring more than about 5 GB. These large data bases are

now commonly referred to as “data warehouses” or “data marts.” The SAS Institute defines a data warehouse as a “separate data store in which the data is stored in a format suitable for business intelligence and decision support systems, in which these systems don’t interfere with the performance requirements of operational systems” (1,2). A data mart is a scaled-down version of a data warehouse, and typically contains between 5 and 15 GB of data, whereas a data warehouse typically contains more than 30 GB of data. Several software developers are currently marketing data warehouse products that help to manage data marts or data warehouses.

Data Base Construction

Because of the potential for very large ITS data bases (greater than 30 GB), most traditional desktop computer-based spreadsheet or data base applications cannot be used to build a data warehouse. Questions of speed, transaction processing time, machine resources and many others must all be considered when determining the system components for the physical construction of the data base. In addition, the relationship of the data base structure itself is of critical importance to processing issues. A large flat file data base, while easy to construct, is poorly suited for the task as there are no shortcuts to value lookups. A more complex relational structure with index tables and shorter lookup paths is far more efficient in processing any query requests.

Access to Data

Most relational data bases require knowledge of a special programming language (e.g., SQL, or structured query language). Ideally, the data warehouse should be accessible to anyone with a desktop computer, without requiring knowledge of programming languages. Several software developers have created access interfaces to data warehouses that do not require a specialized programming language. A simple, easy-to-use access interface to a data mart or warehouse enables a wide variety of users to access and analyze the ITS data.

An advantage of using a vendor created interface package is some minimized project development time to get queries and other analysis routines running on the ITS data. A drawback however is the requirement for the software and the associated costs of the package.

An alternative solution is to construct a completely open standards interface into the data base. While this typically takes more development time, the payoff is in reduced or zero costs associated with accessing and analyzing the data.

Data Versus Information

The sheer size of ITS data bases may make it difficult to convert gigabytes of data to smaller, easy-to-understand information. In many planning applications with ITS data, users may be trying to find patterns or trends over several months or years for the entire freeway system within an urban area. The challenges with large data marts or warehouses are finding and analyzing the appropriate data, then being able to summarize megabytes or gigabytes of data into one useful page of information for

managers or decision-makers. This process of finding and summarizing useful information from large data bases is referred to as “data mining” (i.e., a useful piece of information is analogous to a small nugget of precious metal).

In a popular computer periodical, Gray elaborates on the future of data management (3):

“Perhaps the most challenging problem is understanding the data. There is little question that most data will be on-line--because it is both inexpensive and convenient to store it in computers. Organizing these huge data archives so that people can easily find the information they need is the real challenge. Finding patterns, trends, anomalies, and relevant information from a large database is one of the most exciting areas of data management.”

Privacy

Many privacy advocates are concerned that the advent of the information age means a loss of personal freedom. These advocates are concerned that large data bases containing personal information are kept by government agencies, and that the potential exists for individuals or agencies to merge information from different data bases and construct a detailed account of their personal information and daily activities or habits, including travel activities. Most advocates are not paranoid enough to think that agencies or individuals are currently misusing existing data bases, only that the potential exists for the data bases to be misused.

Privacy issues are an important consideration when information is collected about individual vehicles or persons. These tend to be either individual vehicles (e.g. probe vehicles with AVI transponders or GPS receivers) or persons (e.g. smart cards). In the past typically, information on an individual level was gathered in relation to toll road operations or perhaps on specific studies where drivers have agreed to be probe vehicles or test subjects for data collection. In the future, people may not be aware that vehicle tracking information is being utilized. In all of these cases, information about individual vehicles or persons should either be stored in an anonymous manner (i.e., no way to connect an identification number with a person or vehicle) or not stored at all.

In the current case of TransGuide loop detector data, privacy is not an issue because no information is collected about individual vehicles or persons. However, the Model Deployment Initiative (MDI) in San Antonio includes the deployment of AVI transponders to gather link travel times from participating commuters. The travel times collected by these probe vehicles are a potentially rich source of information to transportation planners; however, storage of the data should make the individual probe vehicle transactions anonymous or aggregate the data. Efforts to incorporate this vehicle-based information into the current ITS data management system for San Antonio should protect the privacy of commuters with AVI transponders.

ITS DATALINK SYSTEM FEATURES

With the above goals and issues in mind, the research team began the design of a prototype ITS data management system that would utilize TransGuide loop detector data. There were several desirable features or attributes for the prototype version of the ITS DataLink system:

- Ability to store and manage large amounts of data;
- Ability to access the data base from remote locations without burdensome or costly software requirements;
- A user-friendly, point-and-click query interface that does not require knowledge of special programming languages or relational data base applications;
- Ability to aggregate and summarize data in point, section, and corridor / facility scenarios;
- Ability to calculate and summarize a given set of performance measures and;
- Ability to output results in a number of different tabular and graphical formats.

The following sections elaborate on these desirable features of the ITS DataLink system.

Data Storage

The researchers agreed that the initial design of the ITS DataLink system should be capable of storing one full year of TransGuide loop detector data. Based on the intended uses of the ITS DataLink system, a decision was made to store the data into five-minute intervals (aggregated from twenty-second periods from the raw data). These design decisions, in addition to the need for additional data base application and temporary file storage space, prompted the research team to select a design specification of 20 Gigabytes for the ITS DataLink system. At this time, the cost of computer storage has become remarkably inexpensive (\$100-200 per GB in 1997 dollars depending on the machine platform).

Data Base Access

It was desirable to have access to the ITS DataLink system from any location without requiring proprietary software. Many data bases require special software to access the system and have licensing requirements for each user. The research team decided to provide access to ITS DataLink using standard Internet protocols (TCP/IP), so that anyone with Internet access could use the system. Password login procedures were deemed to be desirable for the prototype data base so that the privilege levels can be established.

User Interface

The ITS DataLink user interface is very important for enabling novice data base users to perform a wide range of data queries, from very simple to complex. The research team decided that data base users should not be required to know programming or query languages, and that the interface should be a fairly simple, easy-to-use, point-and-click interface. A World Wide Web (WWW or simply web) browser was selected as the interface to ITS DataLink because of its popularity and simplicity of use. Web browsers are increasingly being used to perform complex tasks over the Internet, and even novice computer users generally have experience with using a web browser. Also, Microsoft offers web browser software free of charge on their web site (<http://www.microsoft.com>), and Netscape offers web browser software free of charge to educational and non-profit users (<http://home.netscape.com>).

This Internet is increasingly being recognized as a legitimate tool and location for conducting business. The fastest growing segment of Internet applications is to “web-enable” relational data bases which merely means allowing query and / or update capabilities via a World-Wide-Web browser interface to the database .

There are thousands of examples of web enabled databases on Internet. Some of the more well known transportation examples are the real-time traffic maps for many cities around the nation. These systems are in reality accessing a database of information and displaying it via a graphical summary. Many more examples could be listed but the volume of information makes doing so outside the scope of this paper.

Data Aggregation and Summarization

As described earlier in discussions related to data storage, the loop detector data are stored in the ITS DataLink system at a base level of five-minute time periods. The research team agreed that it was desirable to be able to aggregate the data into a number of different time periods for summarization purposes. Several aggregation time periods were selected for the query interface, including:

- 5 minutes;
- 10 minutes;
- 15 minutes;
- 20 minutes;
- 30 minutes; and
- 60 minutes.

In addition to aggregating the data over various time periods, truly useful analyses depend on aggregating the data across spatial breakdowns as well. As output from the loop detectors in the system, the data represents the smallest spatial resolution--that of a single point on a single lane on the system. To meet the data requirements of multiple applications or uses of the data, the data base had to have the ability to aggregate to other spatial resolutions, including:

- cross lane;
- corridor; and
- facility.

The provision of this ability provides the data base with the power to serve a number of different purposes, from real-time point analysis to corridor operations and even to long-term planning and congestion applications.

Calculation of Performance Measures

The ITS DataLink system was designed to calculate and summarize the performance measures shown in Table 2. The table indicates that performance measure can be calculated for three different spatial levels: point, section, and corridor/system. Because the loop detector stations only provides point data, assumptions about homogenous traffic conditions were necessary to convert the point data to section, corridor, and system measures. These assumptions are usually reasonable in free-flow traffic conditions, but are questionable for congested or stop-and-go traffic.

Tabular and Graphical Output Formats

Any analysis system should be capable of providing performance measure summaries in a number of different formats that are easily interpreted and understood. Tabular formats are fairly common in data summaries, and should be made available through this system. Graphical outputs, though, can better illustrate relationships and trends for large volumes of data. It was also desirable to output data in an intermediate format that could be imported into traditional desktop computer-based spreadsheet or data base applications. The following output formats were selected as being desirable to include in the prototype ITS DataLink system:

- Tabular or columnar data (viewed in web browser);
- Comma-delimited ASCII text (can be imported into spreadsheet or other data base); and,
- Graphical charts in postscript or portable document file (PDF) formats.

ITS DATALINK SYSTEM ARCHITECTURE

Once the feature set was complete, the research team designed the base specifications of the prototype ITS DataLink system. Table 3 summarizes the various components of the system and their specifications. In terms of the software used in development, the ITS DataLink system consists of the five components listed below:

1. Oracle data base management system (DBMS);
2. Common Gateway Interface (CGI);
3. Gnuplot graphics software;
4. E-mail service; and
5. Apache web server.

The Oracle DBMS or relational data base is the root component of the system where the loop detector data are stored. The database also provides some base statistical and aggregation functions. These functions are supplemented by external programs or scripts that are accessed via CGI. CGI, which stands for Common Gateway Interface, is an Internet method of calling one program from within another. This is an important capability of the system as it greatly increases the aggregation, statistical, and data manipulation capabilities beyond what is available in the base Oracle software. The Gnuplot graphics software, which runs on various Unix computer platforms, is a public domain software which enables the system to produce 1-D, 2-D and 3-D graphics. The integration of electronic mail service into the site allows users to mail query results to any valid Internet e-mail address and also allows users of the site to provide feedback to the researchers. The entire package is operated through the Apache web server, a public domain Unix based WWW server software package.

The ITS DataLink system is organized across two Unix workstations from Sun Microsystems. One machine is dedicated to the database, while the other handles the web server functions. In addition to the disk storage requirements of the 5 minute data, the Oracle DBMS occupies about 2 GB of disk space. The external scripts written for the system take approximately GB of storage space. Temporary processing space for the queries requires an additional 4 GB of disk storage space. Figure 1 shows the physical arrangement of the software. To the left of the vertical line represents the server side of the system. The computer represents the web site housed on one workstation, while the data base tables are housed on the second workstation. CGI scripts control the interaction between the various packages. The are to the right of the vertical line in the drawing represents the user's side of the system. The only requirements are a web browser with Internet service and a PDF reader. Both the web browser and PDF reader software are free.

ITS DATALINK SYSTEM COST

The ITS DataLink system is highly scalable. The hardware resources utilized by the research team were Sun Sparc workstations configured with 64 megabytes of RAM and a 2 GB hard disk. Typical cost for this configuration is approximately \$6,000 (1997 dollars). While the prototype system employs two workstations, it is possible to use a single workstation to accommodate all of the system software and loop detector data. Increasing the number of workstations helps to distribute the load across various functions and increases the overall system response speed. In addition to two workstations, two external 9 GB disk drives were utilized to store the loop detector data. External drives have the advantage of being readily transportable from one workstation to another depending on system requirements. A tape backup unit capable of backing up the entire software system software and data completes the physical hardware setup. Total cost for these items of equipment is approximately \$22,000 (1997 dollars)

Any commercial database server, such as Oracle, Sybase, Informix, dB2, which supports standard SQL (Structured Query Language) can support the data base needs. The choice of Oracle for this application was a choice of prior experience to the researchers as well as attractive pricing vs. other available products. With any of the database applications, there is no need for additional programming tools. It should be noted that traditional desktop data bases, such as Microsoft Access, or Corel's Paradox are not recommended for this application due to their limitations in handling very large amounts of data.

CONCLUSIONS

A great deal has been learned from building the ITS DataLink system. Beta testing of the site has proven the concept of building a simple point and click interface to look into large amounts of complex data quickly and efficiently. The performance measures calculated from the data have been used in an operational test to determine bottlenecks on the San Antonio freeway system.

- *Need for a Better Understanding of Data Needs and Uses*--The process of retaining, managing, sharing, and analyzing ITS data for planning and evaluation purposes is not well established among many of the TMCs in the U.S. This research found a major "disconnect" between ITS designers and transportation planners, and that planning data needs were not recognized or understood by many ITS designers. Improved communication and coordination among these two groups should help to address ITS data retention and sharing issues.
- *Plan for Full Breadth of Performance Measures*--The performance measures typically calculated through ITS data only relate to the efficiency and quality of transportation performance. There are several other important components of transportation performance that should be included in evaluating transportation plans or improvements, such as safety, accessibility, and social equity. Although data for these performance

measures are not readily collected by existing ITS components, development of future systems should recognize the data needs necessary to develop the full breadth of transportation performance measures.

- *Query Interface Easy to Use*--The research team found that the ITS DataLink web browser interface made data queries substantially easier to perform for novice computer or data base users. The web browser interface is essentially point-and-click, which most novice users found substantially simpler than similar SQL-based data queries. However, additional input on the query interface would be desirable for future refinements of the system.
- *Map Query Interface Needed*--The current query interface refers to the freeway locations by abbreviated name (e.g., 0U35N) and milepost (e.g., 152.590). Most users are unfamiliar with the TransGuide Phase One system and need supplemental maps and diagrams to interpret queries and data summaries. A map-based interface could improve users' ability to perform and interpret data queries related to specific locations.
- *Corridor and System Analyses Needed*--The prototype ITS DataLink system currently performs queries and presents results for point and link-based performance measures. Tools for aggregating data into corridor and system performance measures should be incorporated into future refinements of the ITS DataLink system. These corridor and system measures provide useful information about regional or sub-regional freeway system performance for managers and decision-makers.

RECOMMENDATIONS

This chapter summarizes the recommendations for this research, which are based upon the major findings and conclusions presented in Chapter 5. The recommendations are offered here as suggestions for future refinements of the ITS DataLink system during the next fiscal year of research activity. The recommendations include the following:

- *Make System Accessible to TxDOT Users in San Antonio*--The research team recommends that the ITS DataLink system be made accessible to interested TxDOT users in the San Antonio District. With a small amount of training and the required passwords, these TxDOT users could have access to the data contained in the ITS DataLink system. The TxDOT users could also provide comments and feedback that can be used to refine the web site and query interfaces.
- *Provide Map-Based Query Interface*--The current prototype of the ITS DataLink system identifies loop detector locations by an abbreviated freeway name and milepost. Most users, however, may not be familiar with the freeway or milepost designations. The research team recommends that a map-based query interface be provided to simplify ITS

DataLink queries. A map-based query interface would allow users to select links or corridors on a plan view of the TransGuide freeway network.

- *Develop Link Aggregation/Dynamic Segmentation Tools*--The current prototype of the ITS DataLink system provides data summaries on a point and link basis. However, the research team has selected numerous performance measures that are corridor or system-based. Link aggregation or dynamic segmentation tools would enable users to summarize data from selected or all freeways into a single performance measure. These corridor and system measures provide useful information about regional or sub-regional freeway system performance for managers and decision-makers.

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the TransLink™ Research Center.

ACKNOWLEDGMENTS

The contents of this paper reflect research conducted for the TransLink™ Research Center located at the Texas Transportation Institute at Texas A&M University. TransLink™ is a public/private partnership between the following partners: Metropolitan Transit Authority of Harris County, Rockwell International, Texas Department of Transportation, Texas Transportation Institute and the U.S. Department of Transportation. The authors would like to thank the TransLink™ Research Center for the opportunity to perform this valuable research.

REFERENCES

1. "Rapid Warehousing Methodology: An Iterative Approach to Building Your Organization's Data Warehouse." White Paper, SAS Institute, Cary, NC, 1996.
2. "The SAS Data Warehouse: Information Systems Management and the Data Warehouse Concept." White Paper, SAS Institute, Cary, NC, 1996.
3. Gray, Jim. "Evolution of Data Management," In *Computer*, October 1996, pp. 38-46.

Table 1. Transportation Management Centers in Telephone Survey

Traffic Operations Center; Phoenix, Arizona

Los Angeles District 7 Traffic Management Center; Los Angeles, California

San Francisco Bay Area Traffic Management Center; Oakland, California

Georgia DOT Advanced Transportation Management System; Atlanta, Georgia

Traffic Systems Center; Oak Park, Illinois

Montgomery County Transportation Management Center; Rockville, Maryland

Michigan ITS Center; Detroit, Michigan

Minnesota Traffic Management Center; Minneapolis, Minnesota

TRANSCOM; Jersey City, New Jersey

INFORM; Long Island, New York

MetroCommute; New York, New York

TranStar Traffic Management Center; Houston, Texas

TransGuide Advanced Traffic Management System; San Antonio, Texas

North Seattle Advanced Traffic Management System; Seattle, Washington

COMPASS; Toronto, Ontario, Canada

Table 2. Performance Measures in ITS Data Management System

Analysis Level	Performance Measures
Point Measures	spot speed (kilometers per hour) person volume or throughput (persons per hour)
Link-Based Measures	travel time (seconds or minutes) average speed (kilometers per hour) person volume or throughput (persons per hour) person-movement speed (person-kilometers per hour) person delay (person-hours)
Corridor or System Measures	average person speed (person-kilometers per hour) total person delay (person-hours) person-miles of travel in congestion (total and percent) person-hours of travel in congestion (total and percent) corridor mobility index roadway congestion index

Table 3. Design Specifications for the ITS DataLink System

Component	Specification
Hardware	Sun Microsystems Ultra1 Sparc Workstations
Operating System	Solaris 2.5 (Unix)
Disk Storage Capacity	20 gigabytes
Data Base	Oracle 7.0
Data Base Access	World Wide Web
Time Aggregation Levels	5 min (base), 10 min, 15 min, 20 min, 30 min, 60 min
Spatial Aggregation Levels	Cross Lane, Corridor, Facility
Available Outputs	Tabular data Comma delimited ASCII text Graphical charts (PDF or Postscript format)

Table Titles

Table 1. Transportation Management Centers in Telephone Survey

Table 2. Performance Measures in ITS Data Management System

Table 3. Design Specifications for the ITS DataLink System

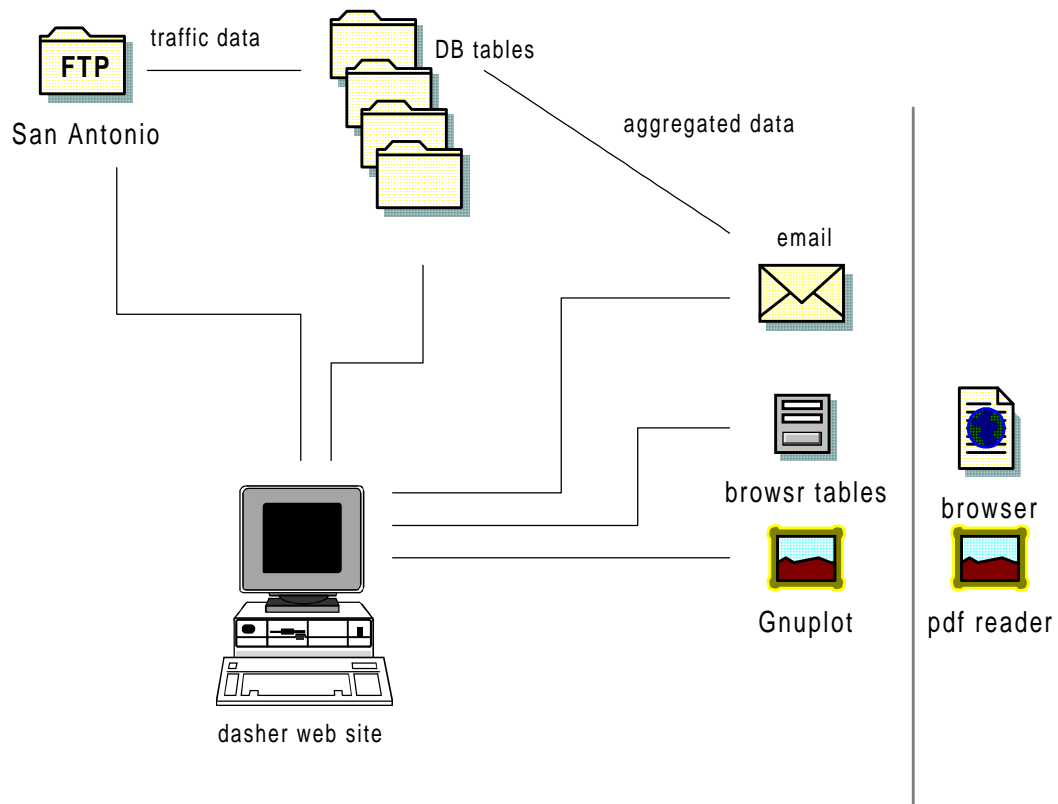


Figure 1. Data Flow for ITS DataLink System